

**UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
WACO DIVISION**

_____)	
Intellectual Ventures I LLC and)	
Intellectual Ventures II LLC,)	Civil Action No. 6:21-cv-00226-ADA
)	
Plaintiff,)	
)	
v.)	
)	
Hewlett Packard Enterprise Company,)	
)	
Defendant.)	
_____)	

PLAINTIFFS' RESPONSIVE CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF AUTHORITIES	ii
I. INTRODUCTION	1
II. LEVEL OF ORDINARY SKILL IN THE ART	2
III. DISPUTED TERMS	3
A. U.S. Patent No. 6,618,736.....	3
1. “intercepting an attempt to [write/read] a data item [to/from] a [storage unit of the first set/a shared storage unit]” (claims 1,6,7,17, 225, 26, 53).....	4
2. “program code for intercepting an attempt to write a data item to a storage unit of the first set” (claim 53).....	6
3. “program code for storing an indication in the first usage map that the Corresponding storage unit of the second set contains valid data” (claim 53).....	9
B. U.S. Patent No. 6,816,464.....	10
1. “route” (claim 1)	11
2. “route statistics” (claim 1)	13
C. U.S. Patent No. 7,783,788.....	14
1. “I/O peripheral subsystem configurations” (claim 1).....	15
D. U.S. Patent No. 8,023,991.....	16
1. “a computer program product recorded on a computer-readable medium, Comprising logic for detecting that a first access point is using a radio Frequency channel ...logic...for instructing the first access point to adjust transmit power” (claim 1)	16
2. “wherein the first access point adjusts transmit power as instructed” (claim 1).....	20
E. U.S. Patent No. 8,725,132.....	23
1. “in order to reduce interference” (claim 1)	23
F. U.S. Reissue Patent No. 42,153	25
1. “sending the poll response communications to the client systems” (claim 1).....	25
G. U.S. Reissue Patent No. 44,818	30
1. “hierarchical token bucket resource allocation” (claims 1, 2, 17, 18, 30, 32, 34)	32

TABLE OF AUTHORITIES

Cases

<i>Huawei Techs. Co. Ltd. V. T-Mobile US, Inc.</i> , No. 2:16-CV-0055-JRG-RSP, 2017 WL 2190103 (E.D. Tex. May 17, 2017)	20, 21, 22
<i>iFly Holdings, LLC v. Indoor Skydiving Germany GmbH</i> , 2015 WL 9258264 (E.D. Tex. Dec. 17, 2015))	30
<i>In re Katz</i> , 639 f.3d 1303 (Fed. Cir. 2015).....	21
<i>Int'l Biomedical, Ltd. v. General Elec. Co.</i> , No. 1:14-CV-397-LY, 2015 WL 7431408 (W.D. Tex. Nov. 20, 2015)	5, 30
<i>IPXL Holdings, LLC v. Amazon.com, Inc.</i> , 430 F.3d 1377 (Fed. Cir. 2005).....	21
<i>MasterMine Software, Inc. v. Microsoft Corp.</i> , 874 F.3d 1307 (Fed. Cir. 2017).....	22
<i>Pacing Techs., LLC v. Garmin Inter. Inc.</i> , 778 F.3d 1021 (Fed. Cir. 2015).....	28, 29
<i>Power Integrations, Inc. v. Fairchild Semiconductor Int'l., Inc.</i> , 904 F.3d 965 (Fed. Cir. 2018).....	30
<i>Regents Of Univ. of Minn. V. AGA Med. Corp.</i> , 717 F.3d 929 (Fed. Cir. 2013).....	29
<i>Rembrandt Data Techs., LP v. AOL, LLC</i> , 641 F.3d 1331 (Fed. Cir. 2011).....	22
<i>Thorner v. Sony Computer Entm't Am. LLC</i> , 669 F.3d 1362 (Fed. Cir. 2012).....	2, 17, 23, 27
<i>Two-Way Media LLC v. AT&T Inc.</i> , No. SA-09-CA-476-OG, 2011 WL 13244905 (W.D. Tex. July 15, 2011)	6, 18
<i>Vitronics Corp. v. Conceptronic, Inc.</i> , 90 F.3d 1576 (Fed. Cir. 1996).....	6, 18
<i>Vivid Techs., Inc. v. Am. Science & Eng'g, Inc.</i> , 200 F.3d 795 (Fed. Cir. 1999).....	24

Statutes & Rules

35. U.S.C.....	2, 7, 21
----------------	----------

Per the Court’s Order Governing Proceedings – Patent Case and Joint Scheduling Order, plaintiffs, Intellectual Ventures I LLC and Intellectual Ventures II LLC (together “IV”), respectfully submit this claim construction brief in response to defendant Hewlett Packard Enterprise Company’s (“HPE”) opening claim construction brief.

I. INTRODUCTION

IV alleges that HPE infringes seven patents owned by IV: U.S. Patent Nos. 6,618,736 (“the ’736 patent”); 6,816,464 (“the ’464 patent”); 7,783,788 (“the ’788 patent”); 8,023,991 (“the ’991 patent”); 8,725,132 (“the ’132 patent”); RE 42,153 (“the ’153 patent”); and RE 44,818 (“the ’818 patent”).¹ IV’s patents teach novel systems and methods in the fields of cloud computing, storage and network virtualization, distributed computing platforms, and wireless networking.

IV’s proposed constructions ask the Court to either construe the disputed terms with their plain and ordinary meaning or to maintain the Court’s constructions from the prior cases where the disputed terms were previously construed.

HPE, on the other hand, asks the Court to read narrowing limitations into the claims. To achieve its goals, HPE takes the contradictory positions that (a) its proposed constructions are correct because they reflect the terms’ plain and ordinary meanings while (b) simultaneously alleging that IV, in also proposing plain and ordinary meanings, is trying to “hide the ball”. *See* Defendant Hewlett Packard Enterprise Company’s Opening Claim Construction Brief, Dkt. No. 31, (“Opening Brief”), at 1. A closer look at HPE’s proposed constructions reveals the purpose behind this seeming contradiction—HPE’s proposed constructions simply do not reflect the plain and ordinary meaning for the disputed terms at all. Instead, HPE labels its proposed constructions as “plain and ordinary,” in one breath, only to then use qualifiers such as “*i.e.*” in

¹ References in this brief to “Ex. __” are to exhibits attached to the Declaration of Jonathan R. DeBlois (“DeBlois Decl.”).

its next breath, to in fact propose further limitations to the terms that would render them as neither plain nor ordinary in nature.

Claims are given their plain and ordinary meaning with exceptions being “1) when a patentee sets out a definition and acts as his own lexicographer, or 2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *See, Thorner v. Sony Computer Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012).

Since HPE has not established that either exception applies, IV agrees with HPE that the terms should be given their plain and ordinary meaning, but without HPE’s unwarranted “clarifying” limitations.

Aside from ordinary meaning disputes, HPE also asks the Court to revisit previously addressed terms from the ’818 and ’464 patents, which this Court has already analyzed and construed. HPE’s proposed claim constructions for these terms are no different from those considered by this Court in the two prior cases. For each of these terms, HPE has not advanced any new rationale for reaching a different conclusion from the one previously adopted by this Court. In fact, HPE does not even *acknowledge* the existence of, let alone address the substance of, the Court’s earlier constructions.

Finally, HPE makes two arguments under 35 U.S.C. § 112. First, HPE proposes a couple of constructions that each fail to include clearly relevant structure from the specification. Second, HPE advances an indefiniteness argument that is at odds with governing law. In both cases, HPE’s arguments should be rejected.

II. LEVEL OF ORDINARY SKILL IN THE ART

The level of ordinary skill in the art does not appear to be substantially disputed by the parties in this proceeding. IV contends that, with respect to all asserted patents, one of ordinary skill in the art is one having at least a bachelor’s degree in electrical engineering, computer

engineering, computer science, or a related field and one to two years of experience in the design or development of networking systems, or the equivalent. Additional graduate education could substitute for professional experience or significant experience in the field could substitute for formal education.

III. DISPUTED TERMS

A. U.S. Patent No. 6,618,736

The '736 patent describes improved systems and methods for creating, managing, and archiving file systems used by virtual private servers, many of which may be running on the same physical host, and therefore typically compete for host resources including file system use and access. *See* '736 patent at 1:26-38, 1:42-45. Such systems are often deployed in multi-tenant environments and need to guarantee that data and processes belonging to one tenant are appropriately segregated from those of other tenants. *Id.* at 1:34-38. This is particularly difficult in the context of a virtualized environment because each virtual private server is a collection of processes that are difficult to isolate from other collections associated with another virtual private server. *Id.* at 1:46-57.

Providing such isolation by giving each virtual private server its own separate physical device for storing files, or its own separate file system, are not viable alternatives. These solutions waste storage space since many different virtual servers make use of the same data and require extensive copying because the filesystems of virtual servers are constantly being updated with new data.

The inventions of the '736 patent solve these and related problems by providing separate file systems to a plurality of virtual private servers even when the virtual servers are utilizing the same or a relatively small number of overlapping physical hosts. *Id.* at 2:42-48. To ensure isolation between individual virtual private servers, without wasting storage space on duplicated

files, and while minimizing the amount of copying required to keep virtual private servers up to date, each virtual private server's file system includes shared and private storage units correlated with one another via a usage map. *Id.* at 2:42-59. The shared storage units contain data that can be used to build a virtual server's initial file system, such as standard application programs, utilities, databases, etc., common to many different virtual private servers, while the respective private storage units contain changes made to the data of one or more of those virtual private servers—data that tends to be private and therefore needs to be siloed from that of other virtual private servers. *Id.* at 8:8-20. The processes that collectively comprise a particular virtual server are tracked, associated together, and linked to that server's private storage area to ensure data integrity. *Id.* at 3:55-4:6. Thereafter, any request for data residing on the shared storage areas, such as a read or write request, is intercepted and redirected to the private storage area associated with the originator of the request. *Id.* at 2:53-67, 4:30-35, 4:58-67. To improve system efficiency, increase scalability, and maximize physical resource usage, data state and location are tracked and identified through a usage map that notes whether requested data is present in the private storage space and, if so, at what location. *Id.* If the requested data is not located within the private storage space, the invention enables it to be copied to available storage space therein. *Id.*

1. “intercepting an attempt to [write/read] a data item [to/from] a [storage unit of the first set/a shared storage unit]” (claims 1, 6, 7, 17, 25, 26, 53)

HPE's Proposed Construction	IV's Proposed Construction
Plain and ordinary meaning, <i>i.e.</i> , receiving an attempted system call, directed elsewhere, to [write/read] a data item [to/from] a storage unit	Plain and ordinary meaning

HPE's proposed construction of the “intercepting” term violates two fundamental claim construction tenets by (1) reading a limitation from a preferred embodiment into the term; and

(2) excluding another preferred embodiment by virtue of reading in said limitation. Opening Brief at 3-4.² HPE disguises this bid by saying the term should be given its plain and ordinary meaning, but then adds a narrowing limitation to the construction, apparently taken in part from a single disclosed embodiment in the specification. *Id.* at 3.

The first part of HPE’s proposal replaces “intercepting an attempt to [write/read] a data item” with “receiving an attempted system call.” *Id.* HPE argues that this must be the case because the specification exclusively discloses a read/write operation as being a system call. *Id.* at 4. This argument fails for several reasons. First, “intercepting” does not mean the same thing as “receiving.” Second, while the specification can certainly inform one of skill in the art as to the plain and ordinary meaning of a claim term, without an express definition or disavowal—neither of which are present here—it cannot be used to redefine an otherwise clear and well-known term. *See Int’l Biomedical, Ltd. v. General Elec. Co.*, No. 1-14-CV-397-LY, 2015 WL 7431408, at *3 (W.D. Tex. Nov. 20, 2015). The term “intercepting an attempt to [write/read] a data item” is clear and unambiguous, and the record contains no definitional statement limiting it to a particular embodiment. The plain meaning is clear from the words themselves—an attempt to write or read a data item is intercepted.

Third, the specification does not exclusively describe the claimed read/write attempt as a system call. A system call is just an example used in the preferred embodiment. *See, e.g.*, ’736 patent at 6:1-2 (“**in one embodiment**, only the system calls 115 that create child processes need be intercepted”), 9:4-6 (“**in one implementation**, a system call wrapper 111 . . . intercepts a system call 115 for writing a data item”) (emphasis added). In fact, the specification expressly

² Pin citations to “Opening Brief” refer to the document page numbers (not ECF page numbers) unless otherwise noted.

discloses alternate embodiments that HPE's proposed construction reads right out of the claim scope. Specifically, the patent discloses an alternate embodiment where, rather than the interception of system calls being used to virtualize a file system, the virtualization is implemented at a lower layer, such as at the virtual disk device level. *Id.* at 12:15-19. The specification further explains that "in such an embodiment, virtualization may occur *once the file system code has turned a system call request into requests for particular blocks on the virtual block device.*" *Id.* at 12:19-22 (emphasis added). Because in this embodiment the claimed intercepting would not be of system calls at all, reading the phrase into the disputed term limits the scope of the claim in a manner not intended by the patentee nor supported by the intrinsic record. *See, e.g., Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1583-84 (Fed. Cir. 1996) (noting interpretation excluding preferred embodiment "rarely, if ever, correct"); *Two-Way Media LLC v. AT&T Inc.*, No. SA-09-CA-476-OG, 2011 WL 13244905, at *5 (W.D. Tex. July 15, 2011) (rejecting defendant's construction as improperly excluding preferred embodiment).

As discussed above, the claim language is clear, unambiguous, and has a plain and ordinary meaning readily understood by even lay people. HPE is simply trying to redefine the disputed term under the guise of allegedly applying a plain meaning, injecting confusion and uncertainty into the claim. The Court should reject HPE's proposal and instead give the term its plain and ordinary meaning.

2. "program code for intercepting an attempt to write a data item to a storage unit of the first set" (claim 53)

HPE's Proposed Construction	IV's Proposed Construction
Subject to § 112 ¶ 6. <u>Function:</u> Intercepting an attempt to write a data item to a storage unit of the first set	Subject to § 112 ¶ 6. <u>Function:</u> Intercepting an attempt to write a data item to a storage unit of the first set

<u>Structure:</u> System call wrapper 111 as shown in Figure 4 and description at 5:17-27, 9:4-7.	<u>Structure:</u> System call wrapper 111, system call vector table 113, Figure 1, Figure 4, 5:17-27, 5:28-39, 5:40-49, 5:50-57, 5:58-67, 9:4-7.
--	---

Since the parties agree that the “program code for intercepting,” term from claim 53 is means-plus-function subject to 35 U.S.C. § 112 ¶ 6, and also agree on the recited function, the dispute here is limited to the corresponding structure. Opening Brief at 6-7. HPE’s proposal for this term excludes structure from the specification that clearly corresponds to the agreed-upon function. *Id.* More specifically, HPE disagrees that system call vector table 113, Figure 1, and several of IV’s proposed specification excerpts provide corresponding structure for the function of “intercepting an attempt to write a data item to a storage unit of the first set.” HPE’s position is incorrect, even as evidenced by its own arguments.

While the entirety of HPE’s objection is baseless perhaps most surprising is its position that Figure 1 should be excluded as structure even though it depicts system call wrapper 111, which *both parties agree* is corresponding structure for the claimed function. *See* Opening Brief at 7. In fact, Figure 1 plays such a prominent role in describing the structure for the agreed-upon function that in arguing that system call wrapper 111 is the only structure disclosed for this term, HPE itself cites to a description of the call wrapper found in the context of Figure 1. *See, e.g.,* Opening Brief at 7 (“[t]he patent confirms that a system call wrapper 111, as described with reference to Figure 1, intercepts a system call 115.” (emphasis added) (internal quotations omitted)). This is because despite appearing graphically in Figure 4, **all** the disclosure with respect to system call wrapper 111 is described in the context of Figure 1. *Compare* ’736 patent at 4:53-54 (noting “As shown in FIG. 1”) *with id.* at 5:20-22, 5:24-25, 5:44-51, 5:58-67, 6:1-5, 6:5-9, 6:12-14, 6:20-21 (all nested under description of Figure 1). Even the brief mention of system call wrapper 111’s presence in Figure 4 refers the reader back to Figure 1 for explanation.

‘736 patent at 9:4-7 (“system call wrapper 111, as described with reference to Figure 1, intercepts system call 115 for writing a data item to a shared storage unit 202” (emphasis added)). HPE’s assertion that Figure 1 relates to a “different embodiment than the one claimed in claim 53” because “[c]laim 53 recites a first usage map but Figure 1 does not depict a usage map” is a red herring. Opening Brief at 8 (internal quotations omitted). While Figure 1 does not expressly show a usage map, it clearly illustrates system call wrapper 111 and is therefore appropriate corresponding structure for the disputed term.³

HPE also disagrees that system call vector table 113 and the passages at 5:28-39 and 5:50-57 provide corresponding structure for the “program code for intercepting . . .” term.⁴ HPE incorrectly argues that the structure described therein plays no part in performing the claimed function. Opening Brief at 8. System call vector table 113 and system call wrapper 111 are inextricably intertwined and are both required to perform the function of intercepting an attempt to write a data item to a storage unit of the first set, as recited in claim 53. For example, for purposes of claim 53, a write attempt is identified as a system call, such as system call 115. *Id.* at 5:17-20. The specification discloses, in one embodiment, that to intercept a system call, a system call wrapper 111 is embedded in the operating system kernel. *Id.* at 5:20-23. But, to intercept the correct system call, the system call wrapper needs to know the exact address of the

³ The veracity of HPE’s position regarding Figure 1 is further called into question considering its heavy reliance on Figure 1 depicting system call wrapper 111 while arguing for its incorrect plain and ordinary meaning interpretation of the “intercepting” term discussed *supra*. That reliance included copying the figure itself into the body of HPE’s brief and outlining system call wrapper 111 in red to call out that it is “used to intercept system calls 115.” Opening Brief at 5.

⁴ It is unclear whether HPE objects to the disclosures at 5:40-49 and 5:58-67 because those two portions of the specification were originally identified as structure by HPE. While those disclosures are missing from the table on page six of HPE’s Opening Brief, the brief only substantively addresses HPE’s opposition to 5:28-39 and 5:50-57. If HPE does object to 5:40-49 and 5:58-67, IV reserves its right to address the same in its sur-reply.

desired system call, which is represented by a pointer and stored in a system call vector table, *i.e.*, “an area in the operating system address space in which addresses of system calls are stored.” *Id.* at 5:28-32. To that end, to implement the actual interception, the appropriate pointers need to be copied, saved, and replaced with pointers to the system call wrapper 111 so that “when a system call 115 to be intercepted is made, the system call wrapper 111 executes instead.” *Id.* at 5:46-48; *see also id.* at Fig. 1. Therefore, to perform the claimed intercepting in the context of means-plus-function claim 53, both the system call wrapper 111 and system call vector table 113 are required. If this were not the case, then a write attempt could never be intercepted, because the system call wrapper would not be able to identify the correct system call address to intercept in the first place. Accordingly, the Court should adopt IV’s proposal, which accurately reflects the disputed term’s corresponding structure, and should reject HPE’s attempt to arbitrarily exclude key structural disclosures.

3. “*program code for storing an indication in the first usage map that the corresponding storage unit of the second set contains valid data*” (claim 53)

HPE’s Proposed Construction	IV’s Proposed Construction
<p>Subject to § 112 ¶ 6.</p> <p><u>Function:</u> Storing an indication in the first usage map that the corresponding storage unit of the second set contains valid data</p> <p><u>Structure:</u> Usage map updating module 404 as shown in Figure 4 and description at 9:18-29.</p>	<p>Subject to § 112 ¶ 6.</p> <p><u>Function:</u> Storing an indication in the first usage map that the corresponding storage unit of the second set contains valid data</p> <p><u>Structure:</u> Usage map updating module 404, usage map 306, Figure 4, 9:18-29, 8:35-45.</p>

Just as with the prior “program code” term the parties disagreement with respect to this term centers around the scope of the corresponding structure. Opening Brief at 8. HPE’s attempt to exclude usage map 306 and the corresponding description of storing an indication in

said map ('736 patent at 8:35-46), as structure for “storing an indication . . .”, is unfounded. Simply speaking, the relevant function is storing an indication in a usage map. HPE takes the untenable position that the corresponding structure for this function should not include the usage map itself. This makes little sense. If the relevant structure for “storing an indication in the first usage map . . .” were not to include the usage map itself, there would be nothing in which to store the claimed “indication.” Furthermore, the disclosure at 8:35-46 includes key structure necessary for the claimed “program code” to perform the stated function, including the programmatic composition of the usage map and indicator. *See* '736 patent at 8:35-46. In any case, the specification makes clear that storing an indication in the first usage map is performed, in part, by the usage map itself. *See, e.g.*, '736 patent at 8:28-57.

B. U.S. Patent No. 6,816,464

The '464 patent relates to systems and methods for route quality testing and prioritization in a communications network comprising multiple interconnected gateways. *See, e.g.*, '464 patent at Abstract, Claim 1. In particular, the '464 patent discloses ways to improve network performance by monitoring, scoring, and evaluating a plurality of candidate routes based on measured route statistics (relating to, *e.g.*, jitter, delay, and dropped traffic) as well as user-selected criteria and routing preferences. *See id.*

For example, the '464 patent discloses embodiments that involve assessing the quality of routes “among various gateways” by “determining the candidate routes to the other gateways, testing those candidate routes, determining candidate route statistics, scoring each candidate route tested, prioritizing each scored route and storing this priority and score information.” *Id.* at Abstract, 2:32-36.

1. “route” (claim 1)

HPE’s Proposed Construction	IV’s Proposed Construction
Plain and ordinary meaning, <i>i.e.</i> , path in the network from the origin of a packet or packets to their destination	Plain and ordinary meaning

HPE’s proposed construction is inconsistent with the plain meaning of the term “route” as it is used in the claims of the ’464 patent. In particular, HPE’s proposal attempts to import the requirement that a route must be a path in the network “*from the origin of a packet or packets to their destination.*” Opening Brief at 10. This additional requirement is inconsistent with the claim language, excludes preferred embodiments described in the specification, and would only serve to add ambiguity to a term already well understood by those of ordinary skill in the art. Accordingly, HPE’s proposed construction should be rejected.

HPE’s sole basis for arguing that a “route” must lead “from the origin of a packet or packets to their destination” is a single, cherry-picked sentence from the specification. *See id.* at 10-11 (citing the ’464 patent specification at 1:25-33). However, a person of ordinary skill in the art would understand that this passage, which appears in the “Background of the Invention” as part of a general overview of existing packet-switching technology, does not relate directly to the “routes” of the ’464 invention, and is not intended to limit the scope of the claims. Specifically, such a person of ordinary skill in the art would understand that the ’464 patent is describing an origin and a destination of a routed packet that does not necessarily correspond to the ultimate origin and ultimate destination of that packet (*i.e.*, much like a route between Waco, TX and Dallas, TX, is not the same thing as a route between 3132 Oliver Street Dallas, TX, and 400 Austin Avenue, Waco, TX).

Further, HPE’s proposed construction is not consistent with the plain meaning of the term “route,” which is used broadly by those of ordinary skill in the art to refer to any path connecting

two devices on a network and is not limited to a path from a packet's origin to its destination. Indeed, in describing preferred embodiments of the invention, the specification consistently uses the term "route" to refer to a network path *from one gateway to another*, and not from a packet's origin to its destination, as HPE's proposed construction would require. For example, the specification explains that "the invention allows for the testing of routes *among various gateways*." '464 patent at Abstract (emphasis added). A preferred embodiment of the invention "provides one or more candidate route(s) *for each gateway from each other gateway*." *Id.* at 7:35-36 (emphasis added). In another embodiment, "routing manager ... determines the latency *between its gateway and the other tested gateways*." *Id.* at 7:50-52.

The use of the term "route" to refer to a network path *between gateways* is also consistent with the explicit language of the claims, which require that each of the "candidate routes" be associated with a "terminating gateway." *Id.*, claims 1, 7, and 13.

HPE's attempt to limit the term "route" to paths that connect a packet's point of origin to its destination would exclude the preferred embodiments described above, and directly contradicts the claim requirement that a route be associated with a "terminating gateway."

Further, as explained in more detail below, HPE's construction directly contradicts the Court's previous construction of the closely related term "route statistics," in which it explicitly equates a "route" with a "network path," and rejects the notion that a "route" refers only to a path "from the origin of a packet or packets to their destination," as HPE argues. *See* Ex. F at IV-HPE-00003804 (construing "route statistics" to mean "Statistics related [to] the quality of a network path (route)").

For these reasons, HPE's proposed construction should be rejected, and the term "route" should be construed according to its plain meaning to refer to any path between two devices on a

network (including, *e.g.*, a path connecting one gateway to another, as explicitly contemplated in the claims and the specification of the '464 patent).

2. “route statistics” (claim 1)

HPE’s Proposed Construction	IV’s Proposed Construction
Plain and ordinary meaning	Statistics related to the quality of a network path (route)

As HPE explains in its opening brief, the term “route statistics” is only in dispute because the parties disagree about the proper construction of the term “route.” Opening Brief at 11. In particular, there is no dispute over the word “statistics,” as HPE admits that “route statistics”—according to the plain and ordinary meaning of that term—are “statistics related to the quality of a route.” *Id.* As explained in the preceding section, the plain and ordinary meaning of the term “route” is “a network path.” Thus, the plain meaning of the term “route statistics” is “statistics related to the quality of a network path (route).”

Indeed, IV’s proposed construction is based on the preliminary construction of “route statistics” that was adopted by this Court last year, in a prior litigation involving the '464 patent. *See* Ex. F at IV-HPE-00003804 (construing “route statistics” to mean “Statistics related [to] the quality of a network path (route)”). Further, the Court’s prior construction of “route statistics” explicitly defines a “route” to be “a network path,” and rejects the notion that a “route” refers only to a path “from the origin of a packet or packets to their destination,” as HPE argues.

For the reasons given above, the Court should adopt the same constructions of “route” and “route statistics” in the instant case as it did before and should reject HPE’s attempt to limit these terms to encompass only paths “from the origin of a packet or packets to their destination.”

C. U.S. Patent No. 7,783,788

The '788 patent is directed to virtualized access to input/output (I/O) subsystems, such as storage or networking subsystems and access to the same. *See* '788 patent at 2:3-8. In particular, the inventions of the '788 patent allow multiple servers to share access to one or more I/O subsystems, logically positioning virtualization logic between the I/O subsystem (*e.g.*, a storage subsystem or a network interface card) and an application server requesting access to that subsystem. *Id.* at 2:8-22. This virtualized access allows storage and network resources to appear as locally attached physical resources from the perspective of an application server, and vice versa. *Id.* at 2:8-17. It further enables a plurality of servers to share a single I/O subsystem such as a storage system or a network interface controller. *Id.* at 2:17-22.

In one non-limiting embodiment, virtual or physical application servers receive configuration information relating to the I/O subsystem sufficient to identify physical I/O subsystem resources available to the operating systems running on the applications servers. *Id.* at 7:1-13. Through the propagation of this information, associations are made between virtual representations of I/O subsystems on one hand, and physical or virtual servers on the other hand. Use by servers of these portions of the resources of the I/O subsystems can be regulated by controlling access to the subsystem at the interfaces leading to/from the subsystem, *e.g.*, a network or storage interface. *Id.* at 7:1-26. By allocating access to physical I/O resources using virtualized representations of those resources and enforcing the resulting virtualized allocations via controlling the interfaces leading to/from the underlying physical resources, the inventions of the '788 patent have enabled a shift away from inefficient static allocations of physical I/O subsystems to application servers. The inventions of the '788 patent have replaced these static inefficient allocations with more flexible allocations of the I/O resources contained within these physical I/O subsystems (or subsets thereof). These more flexible allocations in turn, have led to

increased efficiencies when clusters of servers are deployed with clusters of I/O subsystems, which is one of the key advantages that has driven the widespread adoption of cloud-based computing.

1. “I/O peripheral subsystem configurations” (claim 1)

HPE’s Proposed Construction	IV’s Proposed Construction
Plain and ordinary meaning, <i>i.e.</i> , data indicating the components of the I/O peripheral subsystem	Plain and ordinary meaning.

HPE’s proposal for this term once again attempts to rewrite the claims of the patents-in-suit while hiding behind the fiction that it is merely giving terms their “plain and ordinary meaning.” HPE argues that the disputed term should be limited to “data *indicating the components of the I/O peripheral subsystem.*” *See* Opening Brief at 12 (emphasis added). The patent’s specification makes clear this position is simply incorrect.

In particular, a person of ordinary skill in the art would understand that “configurations” and “components” are not the same thing. This common understanding is reflected in the patent specification, which explicitly uses “components” and “configurations” to refer to different concepts. Indeed, the specification explains that preferred embodiments of the invention may include the same set of *components*, but that these components may be arranged in multiple distinct *configurations*. *See, e.g.*, ‘788 patent at 3:42-46 (“*the application server(s) and the virtual I/O server(s) can be deployed in a variety of configurations*, as illustrated in FIG. 6a thru 6d.”). This passage from the patent’s specification directly contradicts HPE’s proposed construction. If a system with the same components can have multiple different configurations, then a “configuration” *cannot* refer to “data indicating the components” of a system, as HPE argues. Under HPE’s proposed construction, the embodiments described above and shown in Figs. 6a-6d would have the same configuration (because they include the same components).

The specification, however, makes it clear that these figures depict embodiments with “*a variety of configurations*.”

For this reason, HPE’s argument that a list of system components is a “configuration” not only runs contrary to the plain meaning of the term “configuration” as understood by one of skill in the art but is also explicitly disproven by the specification of the ’788 patent itself.

D. U.S. Patent No. 8,023,991

The ’991 patent is directed to a computer program product for adaptively managing wireless communications environments by determining the radio frequency channels and power levels being used by the various access points (“APs”) on the network, determining which of those APs are operating on the same radio frequency channel, and issuing instructions to certain of the APs to adjust their transmit power so as to reduce radio interference caused by their transmissions at other APs using the same channel. ’991 patent at 1:54-64, 5:38-49, 12:59-13:7, 14:40-16:54. In other words, the ’991 patent detects when multiple APs are using the same channel and instructs one or more of them to adjust the power levels of their transmissions to minimize its interference with another “co-channel” AP. *Id.* By using such adaptive power control to reduce co-channel interference, the ’991 patent enables far more APs to operate using the same channels while in close proximity to one another, which in turn, allows Wi-Fi network operators to increase both the number of user devices—or “stations”—operating on their network as well as the amount of bandwidth provided to those stations. *See id.* at 1:54-64, 5:38-49.

1. “a computer program product recorded on a computer-readable medium, comprising logic for detecting that a first access point is using a radio frequency channel . . . logic . . . for instructing the first access point to adjust transmit power” (claim 1)

HPE’s Proposed Construction	IV’s Proposed Construction
-----------------------------	----------------------------

Plain and ordinary meaning, which does <u>not</u> include logic executed at said first access point	Plain and ordinary meaning.
---	-----------------------------

HPE's position with respect to claim 1 of the '991 patent is a sleight of hand. Under the guise of adopting the plain and ordinary meaning of the term, HPE tacks on an extraneous limitation that excludes any embodiment where the claimed “logic” for channel detection and power adjustment is executed at the “first access point,” thus limiting the claim exclusively to “logic” executed by a device that is separate and external to the access point that is being controlled. *See* Opening Brief at 14. HPE's construction finds no support in the claim, excludes preferred embodiments, and artificially narrows the claim scope. The Court should adopt the plain and ordinary meaning of the “computer program product” term *without qualification*, as IV has proposed, and reject HPE's attempt to graft an exclusionary limitation onto the claim.

As noted at the outset, there are only two exceptions to the general rule that a claim should be given its plain and ordinary meaning, and neither is present here. *See Thorner*, 669 F.3d at 1365 (Fed. Cir. 2015). The patentee did not act as his own lexicographer. There is no disclaimer or disavowal of claim scope. Thus, there is no basis for departing from the plain and ordinary meaning for this term. The amendments HPE cites—the only claim amendments made at all during prosecution—were done in response to a basic double-patenting rejection ultimately resolved by a simple terminal disclaimer, and not intended to distinguish the invention from prior art. *See* Ex. C, Pat. Appl. No. 12/827,021 (“the '021 application”), 10/1/2010 Non-Final Rejection; *id.*, 2/1/2011 Amendment at 2; *id.*, 4/12/2011 Terminal Disclaimer. As can be seen from the record—including the sections cited by HPE identified above—the applicant made no substantive remarks or disclaimers that even approach the kind of clear and unmistakable disavowal of claim scope required for the type of narrowing interpretation HPE advances here.

See id. HPE, and not the applicant, is the one imputing a narrow meaning to the claim term that neither the file history nor the law supports.

Simply put, claim 1 of the '991 patent is broad enough to cover an embodiment where the claimed “logic” is executed by the “first access point,” (*e.g.*, by some component or module within the AP itself), as well as an embodiment where the “logic” is executed by some device that is separate from and external to the AP. *See* '991 patent at 42:46-64. The proper interpretation of the claim as having its plain and ordinary meaning—*without HPE’s narrowing qualification*—honors the plain text of the claim, affords the claim its full scope, and finds ample support in the intrinsic record. Specifically, the '991 patent’s specification discloses embodiments where the software for the channel-detection and power-adjustment functionality is executed by and resides *in the AP itself*.

For example, the specification states that the software architecture for the relevant AP optimization functionality of the invention “is advantageously ***implemented in APs*** and STAs in a modular manner for ease of transfer between platforms.” '991 patent at 6:16-19 (emphasis added). HPE’s limiting construction of the disputed term thus cannot be correct, because it would exclude from the scope of the claim the aforementioned embodiment where the software for AP optimization functionality is “implemented ***in APs***.” *See* '991 patent at 6:16-19 (emphasis added); *see also, e.g., Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1583-84 (Fed. Cir. 1996) (noting interpretation excluding preferred embodiment “rarely, if ever, correct”); *Two-Way Media LLC v. AT&T Inc.*, No. SA-09-CA-476-OG, 2011 WL 13244905, at *5 (W.D. Tex. July 15, 2011) (rejecting defendant’s construction as improperly excluding preferred embodiment). Also, for example, the patent explains that “APs and STAs *adjust their power* such that APs and STAs on the same channel can co-exist in an area without

interference,” and that co-channel “*APs continually re-adjust their power level* based on environmental factors ... so that the users’ optimal bandwidth is maintained without undue interference.” ’991 patent at 5:38-49 (emphasis added). The patent clearly discloses embodiments in which an AP adjusts its own power, based on its own logic, and does not limit the claim exclusively to power-adjusting logic executed by a separate device.

Finally, contrary to HPE’s assertion that it is ‘nonsensical and finds no support in the specification’ for an AP to execute logic to detect things about itself or direct itself to take certain actions, the ’991 patent discloses *several* such scenarios throughout the specification. For example, the patent teaches that in one embodiment where a non-invention enabled AP and an invention-enabled AP select the same channel at the same time (channel collision), “the invention-enabled ***AP will direct its own radio*** to the ‘next best channel.’” ’991 patent at 8:11-19. Thus, the invention certainly includes embodiments where some aspect of an AP can direct another aspect of that same AP (here the radio) to take some action. *Id.* Also, for example, the patent describes an embodiment where an AP selects a channel with the weakest signal strength and then “the AP ***12 notes the channel-ID of the channel that it has selected***, the received power level on the channel, and the AP-ID of the AP that generated that power level.” ’991 patent at 8:48-51. Thus, the ’991 patent specifically discloses embodiments where an AP directs itself to do something or detects the channel ID of the channel it itself has already selected.

Moreover, in the latter embodiment just discussed, the AP itself—not a separate device—is also executing the functionality called for by most of the other “logic for ...” elements of claim 1, including detecting channel usage for first and second APs, receiving signals from other APs, and noting power levels of other APs. *See* ’991 patent at 8:48-51 (describing AP detecting channel-ID of its own channel, power level of signals received on same channel from co-channel

AP, and ID of co-channel AP). Accordingly, far from being “nonsensical and find[ing] no support in the specification,” the possibility that an AP itself executes the claimed “logic” makes perfect sense, finds ample support in the specification, and is well within the scope of the claim. Therefore, the Court should give the disputed term its plain and ordinary meaning and should reject HPE's transparent attempt to inject extraneous limitations into the claim.

2. *“wherein the first access point adjusts transmit power as instructed”*
(claim 1)

HPE’s Proposed Construction	IV’s Proposed Construction
Indefinite (mixed method and apparatus claim)	Plain and ordinary meaning.

Claim 1 of the ’991 patent is a straightforward apparatus claim with functional descriptions about the capabilities of components of that apparatus. Claim 1—the sole claim in the patent—is directed to a computer program product comprised of logic for performing various recited functions that involve APs, but the claim does not claim an AP itself, nor does it claim performance of any actions by a user or other person. The final clause in the claim, “wherein the first access point adjusts transmit power as instructed,” simply describes a capability of the first AP with which the claimed computer program logic interacts. The “wherein” clause does not require an action to be performed by the claimed “logic,” but instead defines certain capabilities of the APs that are part of the wireless network environment in which the claimed “logic” operates.

Numerous courts have held that claims with “wherein” clauses reciting functional capabilities, like claim 1 of the ’991 patent, do not render the claim an improper mixed method and apparatus claim. For example, in *Huawei Techs. Co. Ltd. V. T-Mobile US, Inc.*, the disputed claims were directed to a specific structure and its components for use in a particular wireless

network environment. *See* No. 2:16-CV-0055-JRG-RSP, 2017 WL 2190103, at *16-18 (E.D. Tex. May 17, 2017). Despite the claims reciting both the structural component of the apparatus and a “wherein” clause defining actions to be taken by the network, the court held that the disputed claims were not indefinite. *Id.* The court reasoned that the “wherein” clause did not require an action to be performed by the claimed structural component, but rather defined the functional capability of the network environment in which the structural component operated, and thus did not run afoul of 35 U.S.C. § 112. *See id.* In so doing, the *Huawei* court distinguished cases like *IPXL Holdings, LLC v. Amazon.com, Inc.*, 430 F.3d 1377, 1384 (Fed. Cir. 2005) and others, noting that their holdings were premised on claims that called for both an apparatus and use of the apparatus *by a user*. *See Huawei Techs.*, 2017 WL 2190103 at *18. Like the claims at issue in *Huawei*, claim 1 of the ’991 patent is directed to a particular structure—here computer program logic—for performing various functions that involve APs in a wireless network environment, and the inclusion of a “wherein” clause simply defines the capabilities of those APs with which the claimed computer program logic interacts.

In re Katz is similarly distinguishable from the present case. There, the claims at issue covered a system with interface means relating to automated voice messages to individual callers and included a “wherein” clause requiring the performance of specific acts *by certain of the callers*. 639 F.3d 1303, 1318 (Fed. Cir. 2015). The court held the claims indefinite as being directed to both systems and to actions performed by users, thus creating confusion as to when direct infringement occurs. *Id.* In the instant case, claim 1 of the ’991 patent does not recite or otherwise involve a caller, user, or any other person; there is no method step that a person is required to perform; and thus there is no confusion as to when infringement occurs. The Federal Circuit too has distinguished *IPXL* and *Katz* on similar grounds, noting the claims in those cases

required affirmative activities performed by users. *See MasterMine Software, Inc. v. Microsoft Corp.*, 874 F.3d 1307, 1313-16 (Fed. Cir. 2017). Again, claim 1 of the '991 patent contains no such requirement of action by a user, and thus is not indefinite.

Finally, the instant case is distinguishable from *Rembrandt Data Techs., LP v. AOL, LLC*, 641 F.3d 1331 (Fed. Cir. 2011) as well. In *Rembrandt*, while the disputed claim to a data transmitting device did not claim user action, the court held it invalid for indefiniteness because in addition to the four main apparatus elements, the final element recited an explicit method step—"transmitting the trellis encoded frames"—disconnected from the apparatus elements, essentially calling for a method in total isolation from the rest of the claim. *See id.* at 1339. In contrast, the disputed clause of the '991 patent's single claim is preceded by "wherein"—something absent from the method element in the indefinite claim in *Rembrandt*—and merely describes functional capabilities of the first AP with which the claimed computer program logic interacts. *See* '991 patent at 42:46-64.

Like other apparatus claims with functional language found valid by the Federal Circuit, claim 1 of the '991 patent does not reflect an attempt to claim both an apparatus and a method. *See MasterMine*, 874 F.3d at 1315 (cataloguing cases holding apparatus claims with functional language not indefinite). The scope of claim 1 is readily apparent to one of skill in the art, and it is clear when infringement occurs. Just as in *Huawei v. T-Mobile*, the "wherein" clause in claim 1 does not require that the claimed computer program logic perform some method, nor does it require that a user perform some action. It merely specifies the behavior and capabilities of the APs in the network in which the computer program product operates: a network where the access points follow their instructions. As such, claim 1 of the '991 patent is not indefinite.

E. U.S. Patent No. 8,725,132

The '132 patent is directed to an apparatus for (i) controlling and repeatedly adjusting the transmit power level of wireless communication network devices, such as APs, in order to reduce interference, and (ii) transmitting information indicative of a transmit power that devices associated with the AP should use, also so as to avoid interference. '132 patent at 1:53-2:2, 4:67-5:5, 5:39-49, 6:1-9, 6:32-39, 12:62-13:6, 19:36-38, 21:53-57. For example, the '132 patent enables Wi-Fi APs to (i) select an optimal transmit power level, and repeatedly adjust that power level to reduce interference, and (ii) instruct associated end-user devices to transmit at a power level that will also reduce interference. *Id.* This adaptive power control and interference reduction allows Wi-Fi network operators to increase both network capacity and bandwidth provided to end-users.

1. *“in order to reduce interference” (claim 1)*

HPE's Proposed Construction	IV's Proposed Construction
Plain and ordinary meaning, <i>i.e.</i> , for the purpose of reducing interference	Plain and ordinary meaning.

The term “in order to reduce interference,” which appears in claim 1 of the '132 patent, is simple, straightforward, and easily understandable to anyone, including a lay juror. The phrase means exactly what it says and thus should be given its plain and ordinary meaning. HPE does not even attempt to show that either of the two narrow exceptions for deviating from plain and ordinary meaning apply here, because they do not: there is no clear and unmistakable disavowal of claim scope, nor did the patentee act as his own lexicographer. This alone should be fatal to HPE's proposal. *See Thorner v. Sony Comput. Entm't Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2015).

Further, it is not clear what HPE hopes to accomplish by replacing “in order to” with “for the purpose of.” HPE’s opening brief does not explain how, if at all, its proposal would affect the scope of the claims, nor does it identify any ambiguity or confusion that might be resolved by application of its proposed construction. Indeed, HPE seems to argue that “in order to” and “for the purpose of” mean exactly the same thing. *See* Opening Brief at 22. But if this were the case, then there would be no reason for the Court to explicitly construe the disputed term. *See, e.g., Vivid Techs., Inc. v. Am. Science & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed.Cir.1999) (“[O]nly those terms need be construed that are in controversy, and only to the extent necessary to resolve the controversy.”).

However, HPE also argues that its proposed construction somehow relates to narrowing statements made during prosecution. Opening Brief at 21. In particular, HPE puts quotation marks around the word “purpose” to falsely imply that the applicant himself used this term to distinguish his invention from the prior art. This is simply not true: the applicant never used the word “purpose” to define, distinguish or describe the scope of his invention. Instead, as HPE’s own citations show, the applicant’s remarks use the same phrase—“in order to”—that appears in the claim as written. *See* Opening Brief at 22 (citing HPE’s Ex. 9, 8/16/2011 Applicant Remarks); ’132 patent at 43:13-23.

HPE’s arguments about the prosecution history, even though they fail as a factual matter, suggest that the underlying intent of its proposed construction is to narrow the scope of the claims. However, its brief sidesteps the question of exactly *how* its proposal would make the claims narrower or assist the jury in understanding the claims. For now, HPE is just asking the Court to insert “for the purpose of” into the claim. What this phrase actually *means* is apparently

a surprise that HPE is planning to reveal at a later time. The Court should reject HPE's attempt to hide the ball and should give the term "in order to reduce interference" its plain meaning.

F. U.S. Reissue Patent No. 42,153

The '153 patent is directed to methods and systems for the dynamic coordination and control of many network-connected systems (client systems) that collectively perform a distributed computing project, such as the testing of a network site, under the direction of one or more managing systems ("server systems"). '153 patent at Abstract, 2:25-52. Specifically, in the novel approach of the '153 patent, distributed processing of a project collectively occurs on a plurality of client systems that each participate in the project primarily by executing a client agent program. *Id.* at 2:25-52 3:29-46. Throughout such distributed project processing, poll communications from client systems are received by the "server system" to enable it to form a dynamic snapshot of current overall project status. *Id.* at 2:35-39, 20:10-14. The server system then analyzes the dynamic snapshot information to determine if it should decrease or increase the number of client systems actively participating in the project and sends corresponding poll response communications to the client systems. *Id.* at 14:44-56, 18:16-20, 18:50-55. Thus, rather than statically configuring resources or using feedback from client systems solely for resource planning on future projects, the invention of the '153 patent enables dynamic coordination of project activities and resources *while the distributed project is ongoing*, adapting and adjusting to changes in processing requirements as they arise. *Id.*

1. "sending the poll response communications to the client systems" (claim 1)

HPE's Proposed Construction	IV's Proposed Construction
Plain and ordinary meaning, <i>i.e.</i> , sending the poll response communications to each of the client systems	Plain and ordinary meaning.

Once again HPE purports to propose the disputed term’s plain and ordinary meaning, while in effect arguing that the specification and claim language significantly narrow such meaning, without providing any evidence of a clear and unmistakable disavowal of scope or the patentee expressly defining the term. HPE first attempts to recharacterize the claim language to fit its desired narrative by pointing out each instance where the claim itself uses the phrase “client systems.” Opening Brief at 23. HPE reasons that because the phrase appears throughout the claim and also in the disputed term, that “sending the poll response communication to the client systems” must mean sending *to each of the client systems*. *Id.* Except that the claim simply does not say “each of the client systems” nor does it in any way require a poll response communication be sent to “each of the client systems.” HPE’s position is another transparent attempt to redraft the claim.

In any event, HPE’s redraft attempt is based on faulty logic. HPE argues that because the claims do not include a limitation that recites “a subset of the client systems”, they must therefore be referencing “each of the client systems.” *Id.* But the *absence* of the former limitation (“subset . . .”) does not necessarily imply the presence of the latter limitation (“each . . .”). The absence of language such as “subset” merely makes the claim broad enough to encompass both a poll response communication sent to every client system, as well as one sent to just a subset of client systems.

Furthermore, HPE ignores key language in the claim that gives context to the disputed term and additionally illustrates the flaw in HPE’s logic. The immediately preceding step in the claim requires “analyzing the poll communications to determine . . . whether to change how many client systems are active in the at least one project, and if fewer number is desired, including within a polling response communications a reduction in the number of actively

participating clients and if a greater number is desired, adding client systems to active participation . . .” ’153 patent at claim 1. The following step—in which the disputed term appears—claims “sending the poll response communication to the client systems . . . depending upon the one or more decisions reached in the analyzing step.” *Id.*

Following the logical flow and dependency of the claim elements, it is clear that the system first analyzes the project load and the number of client systems that are actively sharing that project load. Subsequently, a determination is made whether to increase or decrease the number of active client systems participating in servicing the project load based on the number of client systems currently active versus the project load at that time. The claim contemplates scenarios where additional client systems are instructed to join active participation and scenarios where currently active client systems are instructed to stop participating. These are alternate scenarios. In the former, a poll response is sent with instructions to join the project to *those client systems determined to be needed*, and in the latter, a poll response with instructions to exit the project is sent *to those clients determined to no longer be needed*. The poll response is indeed sent to multiple client systems; however, there is no requirement to send a poll response to every single client system *within the entire distributed computing platform* when only a fraction thereof may be participating in the project and an even smaller fraction require the instruction within the poll response.

HPE’s argument with respect to the specification allegedly narrowing the disputed term to “each of” is equally flawed because the patentee’s general description of the invention fails to rise to the exacting standard of a clear and unmistakable disavowal of claim scope sufficient to limit the disputed term as HPE proposes. *See Thorner*, 669 F.3d at 1365. HPE’s argument relies on two inapposite citations to the specification—one reciting “dynamic coordination and control

of network connected devices within a distributed processing platform,” and the other noting that “the present invention provides a dynamic coordination and control architecture for network site testing within a distributed processing platform that utilizes a plurality of network connected client devices.” Opening Brief at 23-24. Neither passage comes close to being a clear and unmistakable disavowal of claim scope limiting the disputed term to sending poll response communications to “each of the client systems.” *Id.* In fact, both passages are perfectly consistent with and actually support IV’s position, and the plain text of the claim *as currently written*, that a poll response communication is sent to the client systems, hard stop. That is the plain and ordinary meaning of the term. Nothing in the claim itself nor the file history requires that a poll response communication be sent to ***each of the*** client systems. “The client systems” is not limited to “all client systems” nor is it limited to “less than all client systems”; it is broad enough to encompass either or both. HPE’s attempt to redraft the claim and insert extra words should be rejected.

Additionally, merely using the language “the present invention” is insufficient to meet the exacting standard for a disavowal. HPE cites *Pacing Techs.*,⁵ to support its “present invention” position, but that case is inapposite. In *Pacing Techs.*, the court noted that the patentee used the language “it is a principal object of the present invention to provide . . .,” followed by 19 specific “objects” of the invention. *See* 778 F.3d at 1025. The court stated that this, by itself, is not sufficient to show a disavowal. *Id.* (“the characterization of a feature as ‘an object,’ or ‘another object,’ or even as a ‘principal object,’ will not always rise to the level of disclaimer”). The patentee, however, went much further and “immediately following the enumeration of the different objects of the present invention . . . states that ‘[t]hose [listed] 19

⁵ *Pacing Techs., LLC v. Garmin Inter. Inc.*, 778 F.3d 1021 (Fed. Cir. 2015).

[objects] and other objects and features of the present invention are accomplished, as embodied and fully described herein, by a repetitive motion pacing system” *Id.* (emphasis added, internal quotations omitted). The court reasoned that by stating “as embodied and fully described herein,” the patentee alerted the reader that the invention accomplishes all of its objects and features in the claimed system, and thus made a clear disavowal of scope sufficient to limit the claim language. *Id.*

The holding in *Regents of Univ. of Minn.*,⁶ the case the court in *Pacing Techs.* and HPE quote with respect to the “present invention” language, is similarly distinguishable. The invention in *Regents* was a medical device that used disks to occlude a septal defect. *See* 717 F.3d at 936. The specification described the device as having two separate disks connected via a membrane and included the language “the present invention provides a simple, reliable device for effectively occluding a septal defect. The instant closure device includes [two disks] . . . the membrane of each disk is affixed to [the other disk]” *Id.* The court reasoned that being a physical device that was specifically described as having two physically separate disks in the specification, figures, and prosecution history, it was appropriate to limit the claim scope to that physical configuration. *Id.*

In the ’153 patent, on the other hand, the patentee makes no statements like those in *Pacing Techs.* or *Regents* that the respective courts found were key to triggering the disavowal. The inventions of the ’153 patent relate to a distributed computing platform—something without a static, predefined, exact structure and thus something quite different than a “simple, reliable [medical] device” as found in *Regents*, and without the unmistakable limiting language of *Pacing Techs.* In the instant case, the patentee uses the term “the present invention” 49 times in the

⁶ *Regents of Univ. of Minn. v. AGA Med. Corp.*, 717 F.3d 929 (Fed. Cir. 2013).

specification, many in the context of describing “example embodiments” or optional features that “can,” “may,” or “preferably” be present or that one or more of the inventions “further contemplates.” *See, e.g.*, ’153 patent at 3:46-48, 3:57-62, 4:7-12, 4:13-19. This type of language shows that there are many embodiments and implementations of the inventions disclosed in the ’153 patent, and the discussions of each are illustrative or exemplary, *not exclusionary* and in no way represents an unmistakable disavowal of claim scope. *See Int’l Biomedical, Ltd. v. General Elec. Co.*, 2015 WL 7431408, at *3 (W.D. Tex. Nov. 20, 2015).

Finally, HPE’s proposed construction cannot be correct because it would exclude multiple disclosed embodiments. *See, e.g., Power Integrations, Inc. v. Fairchild Semiconductor Int’l, Inc.*, 904 F.3d 965, 970-72 (Fed. Cir. 2018); *iFly Holdings, LLC v. Indoor Skydiving Germany GmbH*, 2015 WL 9258264, *20-*22 (E.D. Tex. Dec. 17, 2015). Specifically, embodiments are described where each client system can be set to have unique project and poll parameter instructions, such as individualized start/end times, unique poll periods or other individualized criteria based on varying procedures or algorithms. *See* ’153 patent at 17:55-56, 18:38-42. As a practical matter, if every poll response communication—which contain project and poll parameter instructions—were required to be sent to literally each and every client system in the platform each time, it would not be possible to tailor instructions to individual client systems consistent with the above noted embodiment. Therefore, because HPE’s proposed construction is inconsistent with both the law and the facts, the Court should decline to adopt it and instead give this term its plain and ordinary meaning.

G. U.S. Reissue Patent No. 44,818

Like the ’788 patent, the ’818 patent is also directed to systems and methods for the operation of virtualized I/O systems. *See* ’818 patent at Abstract. The ’818 patent, however, focuses on the ability to support a wide range of quality of service (QoS) requirements when

flexibly allocating the resources of I/O subsystems (or subsets thereof) to application servers, by virtualizing the I/O subsystems. *Id.* at 1:66-2:9. This is particularly important because while the virtualization of I/O subsystems greatly increases the flexibility and scalability of computer systems, it also introduce increased management complexity, such as meeting differing QoS requirements of two or more virtual servers, each using logically separate virtual I/O systems that correspond to the same physical I/O subsystem. *Id.* at 1:37-50. A first virtual application server that is running on a physical server for example, might need to use its virtual I/O subsystem for low-latency but low bandwidth transactions (*i.e.*, the low-latency interactions needed by voice over IP servers). At the same time however, a second virtual application server running on the same physical server, might need to use its virtual I/O subsystem for high bandwidth but high-latency interactions (*i.e.*, the high-bandwidth bursty interactions needed by file backup systems), both virtual I/O subsystems being run on the same physical hardware. In such an example, it is challenging to determine how to allocate bandwidth at the single physical interface that might be separating the single physical server from the single physical I/O subsystem in a manner that meets the divergent QoS needs of the two virtual application servers as they transact with their respective virtual I/O subsystems.

Additionally, as QoS policies themselves become more and more granular, it has become increasingly necessary to support numerous classes and levels of access between various configurations of virtual application and I/O servers in a manner that guarantees service but also ensures 24/7 access and data segregation. *Id.*

The inventions of the '818 patent address these and other issues by providing a sophisticated bandwidth allocation scheme across the physical interface that separates a physical application server that hosts several virtual servers, from a physical I/O subsystem that hosts

several virtual I/O servers. The bandwidth allocation scheme is defined by a hierarchical token bucket (HTB), which enables the slicing and dicing (*i.e.*, allocation) of bandwidth to various combinations of virtual application servers and virtual I/O servers using many different types of variables, such as traffic type, or virtual application server identity. *Id.* at 2:10-19. For example, an HTB allocation of bandwidth can force, at one level of its hierarchy, the percentage of overall available physical bandwidth allocated to a particular virtual application server to never drop below a certain threshold. That same HTB allocation can also prevent, at another level of its hierarchy, the bandwidth allocated to a particular type of traffic (*i.e.*, voice traffic) to never exceed a certain absolute threshold. Such a sophisticated multi-level allocation of bandwidth enables a plurality of virtual application servers to each access a logically distinct virtual I/O subsystem, using I/O transactions that are each governed by customized QoS levels of service. *See, e.g., id.* at 3:15-36.

1. “hierarchical token bucket resource allocation” (claims 1, 2, 17, 18, 30, 32, 34)

HPE’s Proposed Construction	IV’s Proposed Construction
Allocation of resources using a class structure arranged in two or more levels, where each class has a bucket of tokens associated with it and scheduling a transmission results in deducting an amount of tokens from a corresponding bucket	Plain and ordinary meaning.

HPE concludes its brief argument in support of the above proposal by stating that its construction “precisely captures each aspect” of the disputed term “while IV’s captures none of them.” Opening Brief at 26. What HPE fails to appreciate, however, is that in *Intellectual Ventures I LLC, et al. v. VMware Inc.*, Case No. 1:19-cv-01075-ADA (“VMware Case” or “IV v. VMware”), Your Honor and technical advisor Dr. Joshua Yi, PhD., considered the same argument from VMware and rejected it, instead construing the term as IV has proposed here—

according to the plain and ordinary meaning. *See* Ex. D (VMware Case Claim Construction Order, Dkt. 103). Despite being aware of this and having the Court’s earlier claim construction order and relevant portions of the hearing transcript, HPE fails to even acknowledge the Court’s previous construction. More surprising though, is that HPE does not point out any alleged errors in the Court’s earlier construction, does not address any new issues or arguments that were not previously before the Court, and does not argue that IV has changed its interpretation or substantive theory in any manner whatsoever. Indeed, as reflected in the below table comparing the construction and support presented to the Court in the VMware Case versus the construction and support HPE presents here, the two constructions are substantively identical, and the support is nearly the exact same.

Presented by VMware	Presented by HPE
“plain and ordinary meaning, which is a tiered arrangement of token buckets where scheduling a message or packet for transmission results in deducting an amount of tokens from a corresponding bucket” <u>Support:</u> 10:14-29	“allocation of resources using a class structure arranged in two or more levels, where each class has a bucket of tokens associated with it and scheduling a transmission results in deducting an amount of tokens from a corresponding bucket” <u>Support:</u> 10:2-11, 10:15-29

See Ex. E (December 3, 2020, email confirming modified construction proposal).

The transcript from the *Markman* hearing held in the VMware Case further confirms this point and reiterates that HPE is just re-packaging VMware’s argument without identifying any errors or new evidence. *See generally*, Ex. G (*Markman* transcript in VMware Case). In explaining to the Court its “modified” proposal reflected in the table above, Counsel for VMware stated that “so if you look at Column 10 of the patent, and the definition really falls in lines 14 through 29, and this is – explains a couple of things. The first sentence says, ‘Hierarchical token bucket can be considered a class based scheduling mechanism.’ . . . the next sentence says ‘HTB includes hierarchical classes where three class types exist: Root, non-leaf and leaf.’” Ex. G at 42.

This is verbatim the same material that HPE quotes in the first paragraph of its argument in support of its proposed construction. Opening Brief at 25. Turning back to the transcript from the VMware Case, Counsel for VMware goes on to recite column 10:21-29 which relate to each class having a bucket of tokens, etc., exactly as HPE does in this case. *Compare*, Opening Brief at 25, *with* Ex. G at 43. HPE concludes that because the specification contains no other disclosure of the “properties of the HTB” the disputed term must be defined as it proposes. Opening Brief at 26. Similarly, in the VMware Case, VMware’s counsel asserted that those citations were “really providing the definition that I think matters for this case.” Ex. G at 43.

As IV explained when the Court first considered this argument less than a year ago, however, those statements are not definitional. Furthermore, the term has a plain and ordinary meaning that is captured by the words of the disputed term itself, and any attempt to further define it interjects confusion and ambiguity into the otherwise clear term or narrows the term in a manner inconsistent with the plain meaning and scope of the claims. *Id.* at 46-47. More specifically, the description at 10:14-29 represents a general overview of the concept of hierarchical token bucket as used in the context of an exemplary embodiment of the “first hierarchical QoS manager,” not the patentee acting as its own lexicographer. *See* ‘818 patent at 9:50 (titling section in question “A.2.2 First Hierarchical QoS Manager”). Perhaps more importantly though, applying HPE’s proposed construction imports significant limitations and confusion into some of the claims. For instance, using claim 32 as an example and applying HPE’s proposed construction, the claim is altered to require a class structure with two or more classes with buckets of tokens and where scheduling a transmission deducts an amount of tokens from a corresponding bucket. Although much of that phraseology is taken from the specification, it also includes many limitations that are otherwise not present in the claim, most

notably the requirements of some type of scheduling and resulting deduction occurring. Without a clear and unambiguous disclaimer or a clear definitional statement of the disputed term, reading in limitations from the specification in the manner that HPE's construction proposes is simply not allowed.

H. Conclusion

Therefore, for the reasons stated above IV respectfully asks the Court to reject HPE's improper proposed constructions and adopt IV's proposals instead.

October 27, 2021

Respectfully submitted,

/s/ Derek Gilliland

Matthew Vella
mvella@princelobel.com
Robert R. Gilman
rgilman@princelobel.com
Jonathan DeBlois
jdeblois@princelobel.com
PRINCE LOBEL TYE LLP
One International Place, Suite 3700
Boston, MA 02110
Tel: (617) 456-8000
Fax: (617) 456-8100

Derek Gilliland
derek@soreylaw.com
SOREY & GILLILAND LLP
133 E. Tyler Street
Longview, Texas 75601
903.230.5600 (Telephone)
903.230.5656 (Facsimile)

COUNSEL for PLAINTIFFS

CERTIFICATE OF SERVICE

I hereby certify that a true and correct copy of the above and foregoing document has been delivered to all counsel of record via the Court's CM/ECF service on this 27th day of October, 2021.

/s/ Derek Gilliland